

PATENT APPLICATION

**ANALYTICAL REPORTING ON TOP OF MULTIDIMENSIONAL DATA
MODEL**

Inventor(s):

Yannick Cras, a citizen of France, residing at
134 Bld de Clichy, Paris 75018, France

Henri Viestro, a citizen of France, residing at
17 Rue Alphonse Moguez, Saint Cloud 92210, France

Ricardo Polo-Malouvier, a citizen of France, residing at
4, Rue Pasteur, Levallois-Perret, 92300, France

Assignee:

Business Objects

Entity:

Large

ANALYTICAL REPORTING ON TOP OF MULTIDIMENSIONAL DATA MODEL

5

CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims priority from U.S. Provisional Application No. 60/194232, filed on April 3, 2000, entitled "System for Analyzing Multidimensional Computer Databases." This application is also related to U.S. Patent Applications No. 09/____, filed April 3, 2001, entitled "Report then Query Capability Multidimensional Database Model (Attorney Docket No. 16499P-000120), and 09/____, filed April 3, 2001, entitled "Mapping of an RDBMS Schema onto a Multidimensional Data Model (Attorney Docket No. 16499P-000130). The above applications and their disclosures are incorporated herein by reference for all purposes.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995

BACKGROUND OF THE INVENTION

The present invention relates to information processing and particularly to creating dynamic interactive reports.

Relational database systems are well-known in the prior art, and include tabular structures and structured query languages used for database queries. To aid in an understanding of the present invention, a glossary of terms is included hereinbelow. Relational database technology has made possible an independence between the way data is physically stored and the way it can be handled logically. Relational technology has been widely accepted in business, industry and engineering as valuable for storing data.

25

Database systems may also be implemented in technology termed OLAP (On-Line Analytical Processing). A key feature of OLAP is a multidimensional view of data. These databases expose data not as entities and associations, as does a Relational Database System, but as sets of values organized in a so-called hypercube. This hypercube structure allows for easy access to data along different dimensions. For instance, a multidimensional database can organize numeric data along three dimensions: Product, Geography, Measures. The product dimension reflects the hierarchy of products in the organization, for example, Consumer and Enterprise products being further divided into more precise product categories

30

down to the final product SKU. The Geography dimension will reflect the corporate organization from the entire company to countries to sales district or ZIP codes. Last, the Measures dimension captures more or less detailed sources of income and expenses, such as sales revenue, maintenance revenue (that may be grouped in a common "revenue" category), or operational expenses. These three dimensions define a system of coordinates. For each valid set of coordinates in this system such as, for example, "All Products," "US," "Revenue," the multidimensional database is able to retrieve a single value (numeric in most cases) that represents the aggregated value of the specified measure for the specified dimension coordinates.

Another feature of OLAP is various analysis models such as parameterized static reporting, slicing and dicing with drill down, "what if" analysis and goal seeking models. These analysis models are provided by standard OLAP tools. These tools serve analytical purposes such as allowing users to select which coordinates in which dimension they want to consider, and must often prompt the user interaction with data through a user interface component called a "grid." A grid allows for easy navigation in a multidimensional cube and enables "drill-down" (descending the hierarchy in a dimension and getting a more detailed view of the data).

From a user perspective, a most important feature of database management software is the user interface and reporting capability. Reports generally present results in user friendly formats, such as graphs, tables, crosstabs, or forms. A novel technique for providing a semantic layer for interpreting and interacting with a relational database is disclosed in commonly assigned U.S. Patent No. 5,555,403 entitled "Relational Database Access System Using Semantically Dynamic Objects," which is hereby incorporated by reference for all purposes.

OLAP products generally provide powerful analytical and navigation tools such as drill down/up, rotation, and slicing and dicing data to provide different views of the data. However, the reporting capabilities of these products are not well developed.

GLOSSARY:

The following is a glossary of some of the terms used in these technologies.

Data: Raw facts or values which are physically recorded and which can be extracted and objectively verified

Information: Anything learned from data, the "meaning" of the data.

Value: An amount of worth

5 **Database:** A computerized collection of data

Relational Database: A database in which all data are stored and organized in tables and in which each field containing a datum is equally accessible.

10 **Relational Database Management System (RDBMS):** The software environment supporting a Relational database.

RELATIONAL DATABASE TERMS:

Record: A collection of fields; the basic accessible element of a file.

Field: An elemental entity of a record.

Relation: An orderly arrangement of data in columns and rows; a table.

20 **Attribute:** A column of a relation, a field of a Tuple.

Tuple: A record of a relational database; one line or row of a table or relation.

Relational Structure: Set of Relations definitions

25

BUSINESS OBJECTS SPECIFIC TERMS:

Infocenter: Environment in which the Users, without any technical knowledge, have "self-service" access to the information system, in a manner controlled by MIS staff.

30 **Universe:** An easy-to-understand partial or total representation of the database, designed for a particular application or group of users.

Business Objects: Objects corresponding to concepts close to the user's everyday business. They are defined by the Infocenter Manager and are the basic elements of a Universe.

Context: set of logical equations linking tables (joins) providing a closed graph of table.

OLAP TERMS:

- 5 **Metadata:** Data about data. How the structure and calculation rules are stored, plus, optionally, additional information on data sources, definitions, transformations, quality, date of last update, user privilege information, etc.

10 **Hypercube:** An OLAP product that has all data stored in a single cube that has all dimensions applied to it.

15 **Dimension:** A structural attribute of a cube which is a list of members, all of which are a similar type in the user's perception of the data. A dimension acts an index for identifying values in a multi-dimensional array.

20 **Cell:** A single data point that arises by selecting one member from each dimension of a multi-

25 **Minicube.** A subset of the hypercube with fewer dimensions than the encompassing hypercube.

30 **Multi-Dimensional Array:** A group of cells arranged by the dimensions of the data.

35 **Drill Up/Down:** Drilling up or down is a specific analytical technique whereby the user navigates among levels of data ranging from the most summarized (up) to the most detailed (down).

40 **Slice:** A subset of a multi-dimensional array corresponding to a single value of one or more dimensions not in the subset.

45 **Slice and Dice:** A method of navigating by calling for page displays interactively through the specification of slices via rotations and drill down/up.

BRIEF SUMMARY OF THE INVENTION

According to one aspect of the invention, a method for performing analytical reporting on top of a multidimensional data model enables a dimension object to be associated with a reporting block, such as a table, crosstab, or chart. The reporting block inherits the data definition of the associated dimension object.

According to another aspect of the invention, each reporting replicates the same set of reporting data for each element of the associated dimension object.

According to another aspect of the invention, dimension objects can be nested so that the reporting object inherits the data definition of both dimension objects.

According to another aspect of the invention, the reporting object is synchronized to the dimension object to display reporting data for each member of the dimension.

According to another aspect of the dimension, OLAP-style analysis is enabled through direct interaction with a report format.

According to another aspect of the invention, the GUI enables associated dimension objects to be filtered by selected member data so that report data is only replicated for selected members.

Other features and advantages of the invention will now be apparent in view of the following detailed description and appended drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 depicts an outline view of a graphical user interface (GUI) utilized in a preferred embodiment to enable analytical reporting.

Fig. 2 is a tree depicting a first level of metadata including dimensions and levels.

Fig. 3 is a tree showing the hierarchy for *Customer* dimension in more detail.

Fig. 4 depicts a document which is a container for one or more working areas called reports.

Figs. 5 and 6 are trees depicting two or more sections at the same level (parallel sections) or at sub-levels (nested sections).

Fig. 7 is a detailed view of a portion of the user interface depicting a Section Zone, an Outline Browser zone, and a Document Zone.

Figs. 8-10 depict the use of the user interface to define sections that replicate the same set of reporting data for various levels of a dimension.

Fig. 11 depicts a user interface for enabling OLAP-style data analysis with a report.

Fig. 12 depicts the cursor changed to magnifying glass to allow drill down when the "+" appears inside the magnifying glass.

Fig. 13 depicts contextual menus that can be activated, for example by right clicking a cell, to enable OLAP-style data analysis on report blocks

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

Reference will now be made in detail to the preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings. While the invention will be described in conjunction with the preferred embodiments, it will be understood that they are not intended to limit the invention to those embodiments. On the contrary, the invention is intended to cover alternatives, modifications, and equivalents, which may be included within the spirit and scope of the invention as defined by the appended claims. No assumption is made on how the data is physically stored. It can be stored into a relational or multidimensional database. For instance, if a relational database is presented as a multidimensional data model as described above, the invention provides the same enhanced reporting capabilities. In one preferred embodiment, the present invention is a software implementation using a three-tiered architecture for providing improved access and reporting on top of a relational or multidimensional database.

The present invention can be utilized in a PC (personal computer) environment, such as is manufactured by International Business Machines Corporation, Apple Corporation, and the like. In addition, all interaction is possible using a "mouse" or the equivalent, and it is to be understood that the following detailed description of the various uses of the improved database accessing capability of the present invention could be realized utilizing such a PC-type configuration.

Fig. 1 depicts an outline view of a graphical user interface (GUI) utilized in a preferred embodiment to analytical reporting. Referring now to Fig. 1, an outline browser is displayed in the left window and a report area is displayed in the right window.

Fig. 2 depicts a first level of metadata including universes and dimensions. The metadata objects that define a universe are: Class, Dimension, Detail, Measure and

Member. The details of a universe are described in detail in the above referenced patent. In the present embodiment, a universe is represented as class and objects.

The metadata dimensions objects that define an OLAP dimension are: Dimension, Level, Hierarchy, Member, and Measure. Figs. 2 and 3 depict the hierarchical structure of a dimension, in this example, customers. The *Customer* dimension includes the levels *County*, *State*, *City*, and *Name*. Also, the properties for members in the *Name* level are depicted. Fig. 3 shows the hierarchy for *Customer* dimension in more detail, showing the members of each level. Additionally, the children of the member *USA* are depicted.

Referring back to Fig. 1, reporting objects can be created by dragging and dropping objects from the Outline Browser Zone to the Report zone. The user interface of the present embodiment allows for arbitrary nesting of objects. For instance, it is possible to insert a table object into a matrix object, so that each intersection of lines and columns in the matrix will contain an instance of the dropped table. It is also possible to introduce "nested axes" in a Table by simply dropping two different dimension objects onto a specific place in the table. Only a few drag-and-drop operations are needed to build such complex reports. The order in which objects are combined does not matter in most cases, which leads to a much easier and intuitive report building process. The data outline mode is controlled by two radio buttons, which determine whether the user browses metadata the objects structure mode or values mode.

The organization of a document utilized in a preferred embodiment of the invention will now be described. Fig. 4 depicts a document which is a container for one or more working areas called reports. A section defines a specific area on a report and is composed of zero or more reporting objects. As depicted in Figs. 5 and 6 there can be two or more sections at the same level (parallel sections) or at sub-levels (nested sections).

Reporting objects included in a section can be blocks such as tables, crosstabs, and charts.

Each reporting object has an associated data definition. Data definitions included in the currently described embodiment are: an axis defined as crossjoin of one or more dimensions, a filter which is an object that selects members; and, a break which is defined as one axis or subpart of an axis available in the table or crosstab. Each reporting object inside a document has one associated logical cube defined by all axes and filters associated with the reporting object.

Data definitions of reporting objects interact through inheritance. In this embodiment, the general rule is that when a reporting object is nested within another reporting object then the nested reporting object inherits the data definition associated with

the parent reporting object. Thus, all blocks within a section are synchronized to the data definition of the section.

According to one embodiment of the invention, the user interface of Fig. 1 is utilized to define sections that replicate the same set of reporting data for various elements of a dimension. These sections can be homogeneous (e.g., one section item for each country) or heterogeneous (e.g., one section item for entire U.S. and one for each of the top ten contributing states). As described above, when a section is defined that includes nested blocks then the nested blocks inherit the data definition associated with the section.

Examples of analytical reporting utilizing the user interface of Fig. 1 and a drag and drop protocol will now be described. Fig. 7 is a detailed view of a portion of the user interface depicting a Section Zone, an Outline Browser zone, and a Document Zone.

The use of the user interface to define sections that replicate the same set of reporting data for various elements of dimension is illustrated by Fig. 8-10. In Fig. 8 a block, in this example a chart, already exists in the document zone. Fig. 8 depicts a chart in the structure mode so that only the header and footer are displayed. In the following example the metadata depicted in the Outline Browser Zone of Fig. 1 are used. In this example, referring to the center panel of Fig. 11, the dimension <Year> is dragged into the Section zone to create a new section having the data definition defined by the object year. As described above, the blocks in the section inherit the data definition of the section. Thus, the block in the form of Table 1 will now inherit the data definition of the year dimension.

As depicted in the center panel of Fig. 9, in the Values Mode a different chart is displayed for each value of the year dimension. In this case, a chart showing Country and Sales is produced for each value of the Year dimension, so that the values of Country and Sales displayed in each chart are synchronized with the corresponding values of the members of the Year dimension.

Figs. 10A-C illustrate an example where there is one section for a particular country, e.g., USA, and sections for the top ten contributing states in terms of sales. As depicted in Fig. 10A, a first section is created (step 2) by dragging the value USA from the Country level of the Customer dimension (see Fig. 3). As described above, the tables (T1, T2, ...) will inherit the data definition of the Country level and will be synchronized along that axis. A filter is then created (step 3) for the country by dragging and dropping the member USA on top of the country cell.

A nested section is then created (step 4) by dragging and dropping the State level between the section <Country> tokens. The tables in the nested section will inherit the data definitions of both the filtered Country and State levels.

5 A filter on the state level is then created by, in this embodiment, utilizing the filter creation GUI depicted in Figs. 10B and C. The user right clicks on the State cell and selects the new filter sub-menu (Fig. 10B) which causes the filter definition wizard (depicted in Fig. 10C) to be displayed. Note that objects can be filtered by Constant, by Values, by a Prompt, or from the top and bottom. In this example "Top" and "10" are selected and objects are filtered by sales revenue so that the tables will displayed having values for each of the top
10 ten states in terms of sales revenue.

Another feature of the currently described embodiment is the enablement of direct user interaction with a report to perform the same analysis that would be performed by an OLAP tool. Fig. 11 depicts a user interface for enabling OLAP-style data analysis with a report. The magnifying glass with "+" defines the "Drill and Replace" mode which is the default mode. The expansion icon with the "+" and "-" selects either "Drill and Replace" or "Drill and Expand."

Interaction with cells or blocks is enabled utilizing various GUI techniques. For example, as depicted in Fig. 12, the cursor can be changed to magnifying glass to allow drill down when the "+" appears inside the magnifying glass. When the cursor is over a cell where no drill down is possible, but a drill up is possible, the "+" inside the magnifying
20 changes to a "-".

Alternatively, as is known in the art, contextual menus, as depicted in Fig. 13, can be activated, for example by right clicking a cell, to enable OLAP-style data analysis on report blocks.

25 In a preferred embodiment, the processes described above realized in software, stored on a computer readable medium, executed by a processor. The storage medium can be magnetic or optical storage and includes digital signals encoded in an electromagnetic signal.

The invention has now been described with reference to the preferred embodiments. Alternatives and substitutions will now be apparent to persons of skill in the
30 art. For example, although the embodiment described created sections utilizing a drag-and-drop GUI technique, contextual menus or other GUI techniques can be utilized. Accordingly, it is not intended to limit the invention except as provided by the appended claims.